

IP-ATM Convergence/Internetworking

Umair Hoodbhoy
Internet Protocols
April 1999

ABSTRACT

The internetworking problem is a rather complex one for the networking industry. Since the interface between layers outlives the technology used in each layer and TCP/IP is here to stay, any new technologies must be compatible with it. ATM promises very high-speed transfer of data, voice and video while delivering QOS guarantees. However, ATM is connection-oriented whereas most modern data networking protocols are connectionless. Hence, applying ATM technology to data communication is an intricate problem. SONET/SDH is emerging as a prospective solution to this problem. This paper compares running IP over SONET/SDH with running IP over ATM.

INTRODUCTION

The Internet has made irreversible changes to our lives as it continues to grow at a rapid pace. The number of hosts on the Internet has doubled approximately every year since 1989; as of January 1999, the number was over 43 million [1]. Distributed applications are increasingly bandwidth-hungry, and the Internet is demanding more than what legacy LAN architectures can provide. The demand for providing multimedia (voice, video, and data) over the Internet is growing rapidly as well. The internetwork is the critical tool for information flow. There are calls for internetworks to cost less yet support the emerging

applications and higher number of users with increased performance (beyond Best Effort). A corollary of internetworking is 'universal interconnection'; that is, the user is provided an interface that is independent of the network. The advantage of using TCP/IP is that all sub-networks are treated equally.

Until recently, local and wide-area communications have remained two different issues. Bandwidth is free in LANs whereas it is the predominant cost in WANs. Hence, delay-sensitive traffic such as voice has been treated differently from data. However, the Internet is changing these trends because applications such as voice and real-time video require better, more predictable LAN and WAN performance [2].

In order to provide high-speed data and video capability, yet reduce the cost of voice and data transmission, the Integrated Services Digital Network (ISDN) was standardized. Narrowband ISDN appeared in 1980 and uses circuit-switching technology and 64-kbps channels for switching. It soon became apparent that higher bandwidths would be needed. Broadband ISDN appeared a few years later as a packet-switching technology that supports very high data rates (100s of Mbps) and making Quality of Service (QOS) guarantees. Asynchronous Transfer Mode (ATM) was chosen for this technology in 1988. The Broadband ISDN protocol suite is shown in Figure 1.

Signaling	Connection-Oriented	Connectionless	Variable Bit Rate	Constant Bit Rate
ATM Adaptation Layer				
ATM				
Physical				

Figure 1.

This paper will discuss two ways to integrate IP over the ATM and Physical layers (LAN Emulation and Classical IP over ATM) and then a way to run IP over an alternative network (SONET/SDH).

PROBLEM SPACE DEFINITION

The two most common ways to run IP over ATM are LAN Emulation (LANE) and Classical IP over ATM (CIP).

LANE is an ATM Forum standard that allows existing LAN-based (Ethernet or Token Ring) clients to take advantage of the benefits of ATM (e.g. speed and QOS) without making any hardware or software modifications to the client machines. In this sense it emulates a 'legacy' LAN to form an emulated LAN (ELAN). LANE is flexible in that provides a scalable transition to ATM. Since it is meant for ATM hosts to co-exist with legacy IP hosts (yet use connection-oriented data transfer), it uses the ATM Adaptation Layer 5 protocol (AAL5) in the Broadband-ISDN protocol suite. It acts as a bridging layer between AAL5 and the upper-layer connectionless protocols above it. An ATM switch and an ATM-to-LAN converter would separate each ATM and Ethernet or Token Ring host. LANE performs the bridging between the Media Access Control (MAC) layer on the Ethernet side and the AAL5 layer on the ATM side [3].

Besides the problems to be discussed later, LANE has its own performance cons. It cannot use large MTUs, so IP packet fragmentation occurs frequently. Moreover, its support for broadcast and multicast

is not as efficient as CIP's.

CIP is the Internet Engineering Task Force's (IETF) specification. It divides an ATM network into several logical IP subnetworks (LIS) consisting of end systems or hosts with the same IP prefix. Hosts in the same LIS communicate with each other through end-to-end ATM connections (Virtual Channels) [4]. Since broadcasts are not allowed in ATM, ATMARP servers are used in logical subnets to resolve the target IP addresses into native ATM addresses. The service that provides this is the ATMARP protocol and its scope is limited to the LIS [5]. So each host's ARP cache maintains an IP to ATM address mapping instead of an IP to MAC address mapping. If a host needs to communicate with a host outside its LIS, it goes through an IP router. This was proving to be a bottleneck until the Next Hop Resolution Protocol (NHRP) emerged [6]. NHRP makes it possible to use "short cuts" from a node on one LIS directly to a node on another LIS, thereby avoiding router hops between the LISs. Incidentally, LANE faces a similar problem when inter-ELAN traffic has to pass through a router. So the Multiprotocol Over ATM (MPOA) group was formed by the ATM Forum. MPOA combines LANE and NHRP functionality to support IP routing as well as LAN bridging over an ATM network [7]. It should be noted, however, that NHRP opens a new can of worms and according to RFC 2382, "it is a topic for further study to determine if significant benefit is achieved from short cut routes vs. the extra state required" [8].

Both LANE and CIP are very difficult to implement and manage. Moreover, since ATM is connection-oriented and IP is connectionless, inefficiency can easily creep in. For example, only until the LANE 2.0 specification, QOS was not possible within LANE so only Available Bit Rate (ABR) or Unspecified Bit Rate (UBR) connections could be used [9]. QOS guarantees require a Variable Bit Rate (VBR) connection for each distinct flow [10]. Running TCP/IP over ATM can introduce redundancies as well; ATM already has functionality for control so TCP is useless.

But there is a bigger problem associated with running IP over ATM. This is the infamous 'cell tax' issue [11]. Roughly 10% of raw bandwidth (5 bytes for every 48 bytes payload) is wasted in the ATM cell header. That overhead actually raises to 24% because ATM does not segment small packets well. Unfortunately, about 45% of the most popular packets are 44 or 40 bytes (acknowledgements) which are so small that they cannot be encapsulated within an ATM cell as described in RFC 1483 [12]. This entails to only 135.63 Mbps of ATM payload on a 155.52 Mbps link. Clearly, there is plenty of room for improvement in the field of running IP over ATM.

SOLUTION SPACE EXPLORATION

SONET (Synchronous Optical Network) was designed as a high-speed physical layer for digital optical transmission over optical fiber. It was standardized by ANSI and is the United States version of ITU-T's Synchronous Digital Hierarchy (SDH). By placing the ATM cells in a synchronous time-division envelope, SONET/SDH is one way to transmit ATM cells (the other is to use a continuous stream of cells and hold the receiver responsible for being able to separate cells on the 53-byte boundary [13]). So an early implementation of running IP over ATM was to use a SONET/SDH physical layer as shown in Figure 2 [14].




Figure 2

However, as we have seen, there are several complications with running IP over ATM. So an alternative approach was sought.

There is a bright future for running IP over SONET/SDH using the Point to Point Protocol (PPP) as a link layer protocol. RFC 1619 defines a format for running PPP over SONET/SDH. Using PPP encapsulation provides much higher (25-30%) throughput and costs less than running IP over ATM because of its relatively low overhead. So, the trend is changing such that IP is encapsulated within PPP and transmitted using SONET as shown in figure 3.




Figure 3

SONET is composed of Synchronous Transport Signal (STS) levels which have SDH counterparts known as Synchronous Transport Modules (STM). The lowest STS level is 51.84 Mbps and all SONET trunks run in multiples of this. STM-1, the lowest STM, is 155.52 Mbps, which corresponds to STS-3. There are also Optical Carrier (OC) levels that correspond to the SONET designation. So OC-3 is the same as STS-3 and STM-1 and has a data rate of 155.52 Mbps. This is the basic rate for PPP over SONET/SDH. The actual payload rate is 149.76 Mbps.

SONET frames are basic blocks of 810 bytes sent every 125 μ s [15]. In the case of running PPP over SONET, for the basic STS-3 rate, 3 blocks (2430 bytes) are used. A 9x9 matrix is assigned for transport overhead while the remaining 261 columns are user data, known as Synchronous Payload Envelope (SPE). This is shown in Figure 4.




Figure 4.

PPP frames are encapsulated in a manner like ISO's High-Level Data Link Control (HDLC). However PPP uses byte stuffing compared to HDLC's bit stuffing. RFC 1661 discusses PPP in more detail.

SONET/SDH and PPP work well together because the former, by definition, is a point-to-point circuit, and the latter was designed as a standard means of communicating over point-to-point links [16].

RFC 2171 introduced Multiple Access Protocol over SONET/SDH (MAPOS) as a method for transmitting HDLC frames over SONET/SDH [17]. This has more features than PPP as a link-layer protocol as it supports multiple access, broadcast and multicast and operates over SONET in both the LAN and WAN environments. The mechanism for transmitting packets is almost trivial: the node simply fills in the destination address of an HDLC frame and sends it over a SONET/SDH link. Moreover, since MAPOS is a link layer protocol, IP can run over it independently and IP datagrams can be encapsulated in HDLC frames of MAPOS [18]. Each node on a MAPOS network maintains an "ARP cache" that can be modified manually or by the Address Resolution Protocol (ARP). The ARP cache maps destination IP addresses to their corresponding 8-bit HDLC addresses in a manner similar to Ethernet ARP. RFC 2176 discusses IPv4 over MAPOS in greater detail.

CRITIQUE OF SOLUTIONS/IDENTIFICATION OF OPEN PROBLEMS

The biggest advantage of running IP over SONET/SDH, whether using PPP or MAPOS as a link-layer protocol, is the elimination of ATM cell tax. The other reason why IP over SONET/SDH might be a better, more feasible solution than IP over ATM is that ATM is too intricate with its connection-orientation. Setting up virtual circuits requires a lot of processing power.

However, ATM provides other services that PPP or MAPOS cannot. For example, ATM allows multiple secure Virtual Circuits (VCs) on the same physical interface, provides traffic management, and can support QOS guarantees, which are difficult to achieve in PPP. Also, ATM is a multi-destination protocol while PPP can only reach one destination per link. MAPOS amends this problem although it is still 'under construction'. SONET/SDH has a feature called Automatic Protection Switching that switches the circuit when faults occur or signals are lost. The downside of this is the delay incurred while switching (up to 50 ms). Some SONET/SDH implementations use some overhead bytes in proprietary manner and that, along with other factors, can cause interoperability problems. When running IP over SONET/SDH, all data traversing the link must be IP, which is a drawback because 155 Mbps of IP data is not easy to come by. IP over SONET/SDH does not statistically multiplex as well as ATM. Also, ATM can be run on lower-speed links.

Nevertheless, IP over SONET/SDH has a bright future and many corporations are realizing this. For example, Internet Servers Inc, the company that hosts high-traffic sites such as Yahoo Inc. and Netscape Communications Corporation implements IP over SONET/SDH.

SUMMARY

As the demand for multimedia (voice, video) over IP increases, the networking industry is facing new challenges to provide high speed networking at the LAN and WAN level whilst fulfilling QOS requests. ATM can provide this but it was put to the test to determine whether it could support IP well. Efforts to run IP over ATM have been made and both IETF and the ATM Forum have defined standards. Both

Classical IP and LANE have pros and cons. Routing protocols such as NHRP and MPOA have helped the cause of CIP and LANE respectively. The bottom line, though, is that since ATM is connection-oriented whereas IP is connectionless, implementing IP over ATM is a complex issue. So IP over (PPP over) SONET/SDH was looked into as an alternative. IP over SONET/SDH is less costly and introduces less overhead than when ATM is involved. PPP introduces a few problems but using MAPOS as the link-layer protocol instead can solve them. Management of IP over SONET/SDH is generally difficult but because high-volume point-to-point configurations are easily implemented, it is gaining popularity in corporations.

BIBLIOGRAPHICAL REFERENCES

1. Network Wizards, www.nw.com.
2. *ATM Internetworking*, Anthony Alles; <http://www.cisco.com/warp/public/614/12.html>
3. *ATM LAN Emulation: An Inside Look at Version 1.0 of the LANE Specification*; Bob Klessig; <http://www.3com.com/nsc/500617.html>
4. RFC 2225; *Classical IP and ARP over ATM*; M. Laubach; January 1994.
5. *Methods for Implementing IP on ATM Networks*; Robert E. Larson
http://www.odyssea.com/whats_new/ipoveratm/atm.html
6. *IP over ATM: Classical IP, NHRP, LANE, MPOA, PAR and I-PNNI*; Jun Xu
http://www.cis.ohio-state.edu/~jain/cis788-97/ip_over_atm/
7. RFC 1932; *IP over ATM: A Framework Document*; Cole et. al.; April 1996
8. RFC 2382; *A Framework for Integrated Services and RSVP over ATM*; Berger et. al.; August 1998.
9. *TCP/IP over ATM: A No-Nonsense Internetworking Guide*; Berry Kercheval; Prentice Hall, 1998
10. Cisco Documentation;
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2008.htm>
11. *Cell Taxes: A Necessary Evil?*; Phasecom Inc; <http://www.speed-demon.com/sonet.html>
12. James Manchester's notes; <http://hibp.ecse.rpi.edu/~manchj/>
13. *ISDN and Broadband ISDN with Frame Relay and ATM*, William Stallings; Prentice Hall, 1999
14. Dr. Raj Jain's notes; <http://www.cis.ohio-state.edu/~jain/talks/ipsonet.htm>
15. *Computer Networks*; Andrew Tanenbaum; Prentice Hall, 1996
16. RFC 1619; *PPP over SONET/SDH*; W. Simpson; May 1994
17. RFC 2171; *MAPOS - Multiple Access Protocol over SONET/SDH Version 1*; Murakami et. al.; June 1997.
18. RFC 2176; *IPv4 over MAPOS Version 1*; Murakami et. al.; June 1997.